MARTIN'S PHILOSOPHICAL NOTES

The Social Costs of Worshipping Computers

Martin Cohen

The coronavirus crisis – at last! – begins to recede, but it leaves behind it several big new questions. One of which is: What is the role of computers – or to be precise, sophisticated computer models – in creating the 'crisis'?

Flashback to early 2020 and the first reports of the new virus. It's a matter of public record that when Neil Ferguson's model produced a calculation showing that the United Kingdom's health service would soon be overwhelmed with severe cases of Covid-19, and that Britain would suffer more than 500,000 deaths, Prime Minister Boris Johnson, who up till then seemed to be flirting with notions of 'herd immunity', instead swiftly announced stringent new restrictions on people's movements to 'slow the spread'. The computer, not the politicians, decided the policy. And likewise, when the same Imperial College London model warned[1] that, unless the country locked down soon, the United States would end up with 2.2 million virus deaths.

When policy makers consider the print-outs of computer models like this, quite how the models actually work (it's complicated) is of little interest. But the lack of curiosity, let alone scepticism, is not surprising, given that we now live in a world in which scientists with computers have replaced priests with a crucifix as the sources of unchallengeable truth and wisdom.

Instead, governments, journalists and the general public simply accept the assurance of the researchers that the models are accurate. This is despite the fact that computer modellers have a terrible record of ludicrously misjudging previous non-epidemics.

Infamously, in 2005, Ferguson predicted that up to *200 million* people could be killed from bird flu – a story I recounted in my book, *Paradigm Shift* (2015).[2] His department, which is closely linked to the pharmaceutical industry, recommended that an international stockpile of three million courses of antiviral treatment be built up. This, the computer model indicated, would be enough to contain an outbreak.

In the end, only 282 people died worldwide from the disease between 2003 and 2009. How did Ferguson arrive at his figure? Well, he told the *Guardian* newspaper [3] that 'around 40 million people died in 1918 Spanish flu outbreak…. There are six times more people on the planet now so you could scale it up to around 200 million people, probably.'

That back-of-the-envelope calculation just looks dodgy. However, the bigger problem is that with computer modelling, as with anything stated by a computer, the results, crisply printed out in milliseconds, appear dispassionate and compelling. Yet often, key assumptions programmed into the computer are based on these sorts of lazily arrived at, back-of-an-envelope calculations. This despite the reliability of outputs being essentially determined by the assumptions and data plugged into the models.

Of course, academics and computer modellers insist that such problems are always 'about the past'. The latest model is much better than those old, discredited ones! But you'd have to wonder, as the people building the models haven't changed, even if their computers have got fancier. And when teased out of their grubby offices, computer modellers, just as much as

members of the public, freely admit to 'knowledge gaps', and to having to occasionally rely on what we might call 'fast thinking' to get the models up and running. Put another way, plugging in easily available assumptions – instead of doing expensive and time-consuming research.

The bottom line is that computer modelling is not entirely 'science', and there's a lot of guesswork, too. Maggie Koerth, Laura Bronner and Jasmine Mithani raised many good questions in an article for 538 Magazine, entitled 'Why it's so freaking hard to make a good COVID-19 model'.[4] As they put it there:

> Every variable is dependent on a number of choices and knowledge gaps. And if every individual piece of a model is wobbly, then the model is going to have as much trouble standing on its own as a data journalist who has spent too long on a conference call while socially isolated after work.

Academic point irrelevant to policy? Not at all. Take a closer look at that influential model of Imperial College, run by the aforementioned academic, Neil Ferguson and his team. The underlying structure of the model was supposed to represent the behaviour of the Spanish influenza virus. That's not because Covid was thought to be like influenza, but because that was the model that they had to hand.

One difference is that flu is basically more infectious. Another supposed difference was that Covid (particularly in the early days) was considered to be much more deadly. So, to make the model fit Covid, a few variables were tweaked. Notably the case fatality rate, which went up ten-fold.

This raises the question – do computer modellers have some secret insight into key facts about things like the corona virus that elude medical specialists? Guess what – the answer is 'no'. And so the result was – and is – that computer modelling of the virus, that seems so detailed and comprehensive and feeds fear of the disease, was based on very shaky assumptions.[5]

But take another example from Professor Ferguson's model. In the model, fewer than 5 per cent of people are infected, yet in the real world, the number, according to researchers at Oxford's Evolutionary Ecology of Infectious Disease lab, could be ten times that, maybe even over 50 per cent! These unaccounted cases are people who have coped with the virus without appearing in hospital or other official statistics. They are 'the herd' amongst whom the virus can no longer circulate. But this group all but disappeared with the click of a computer key, say in the Imperial College model. And unfortunately, as logicians say, *any* conclusion can follow – perfectly logically – from false premises.

Ferguson refused to share the workings of his model with other researchers.[6] He wrote:

> I'm conscious that lots of people would like to see and run the pandemic simulation code we are using to model control measures against COVID-19. To explain the background – I wrote the code (thousands of lines of undocumented C) 13+ years ago to model flu pandemics….

He said that instead, he had asked Microsoft to produce a new model based on his one that other people could then pore over.

At other times, though, the garrulous Ferguson has been quite frank about the status of crucial figures, like that for the 'Case Fatality Rate' (CFR). Not here, diligent studies of medical data. Instead, the figure for the virus was simply plucked from the general swirl of misinformation. Which figure to use was essentially a *political* decision. Which is why the already existing, real-life case study of the spread and deadliness of the coronavirus that was the cruise ship, the *Diamond Princess*, was of less interest to them than the purely speculative claims of the World Health Organisation,[7] who early on used a CFR of 3.4

per cent, even if now, grudgingly, they say it is much less.

Let's go back to the *Diamond Princess*. I read about this early on in the virus story, in the context of a very clear piece[8] by the sceptical epidemiologist Professor John Ioannidis for *STAT* magazine. It was in this piece that Ioannidis presciently warned that the emergency measures to combat the virus could cause far more deaths than the virus itself. However, as far as how deadly the virus was, the story tells how several passengers on the cruise ship contacted the virus – turning the whole ship into a kind of giant, awful, experimental test-tube in which the infectiousness and spread could be seen in real time. Indeed, the virus quickly spread through food service workers, particularly those cooking for other members of the crew. Eventually, of the 3,711 passengers and crew aboard, some 700 tested positive, and seven people died. This was rapidly adopted as the benchmark 'case fatality rate' – 1.0 per cent. Ferguson's Imperial model[9] uses 0.9 per cent, or in grander language: 'anyone infected with SARS-CoV-2 has a probability of dying with mean 0.9% with sigma = +/-0.1%'.

Politicians and journalists can easily be persuaded that such things are all very 'scientific', but as Ioannidis says, the crucial fact about the Covid virus, clear right from early on, was its steep age-gradient – that it affects elderly people, and leaves the young relatively unaffected. (Influenza, which the models originally were designed for, is also dangerous to elderly people, but is significantly different in that it is dangerous to children under 5.) Add to which, a cruise-ship sample is not a normal mix of people but, instead, is a largely elderly population, in which the death rate from Covid-19 is bound to be much higher. As Ioannidis pointed out, right from the start of the crisis, 'Projecting the Diamond Princess mortality rate onto the age structure of the U.S. population, the death rate among people infected with Covid-19 would be 0.125%' – only one eighth as high.

Ferguson and his team, however, assumed a uniform fatality rate applied to infected individuals, ignoring the evidence that SARS-CoV-2 has a toll skewed heavily by age and co-morbidities. The suspicion was that Ferguson and his team were inclined to make headlines and money with eye-catching predictions. A model that shows the virus spreading harmlessly through the population before fizzling out would not bring in either media attention or emergency funding.

Perhaps, too, using an inappropriate model might be because it might be very complicated to change a flu model to accurately reflect things like the medical risk and age factors. I'm not sure. 'Knowledge gap!' What I *am* sure about is that the consequence is a huge policy error. Because the computer modellers treated everyone as equally infectious and equally vulnerable, so too did the policies. Likewise, we still don't really know if there is any such thing as 'asymptomatic transmission', or at least how dangerous it might be; and we do know, because it is common sense, that the very old and people with existing illnesses are exponentially more vulnerable to corona viruses (plural), and thus policy should have been centred on that: 'focussed prevention'.[10]

Anyway, the computer modellers, along with other media 'experts', and medical specialists continued plugging in figures as high as 10 per cent for coronavirus fatalities, and failing to distinguish between which kind of person was most at risk. Had they done so, they would have been forced to conclude, as Ioannidis calculated early on, that a reasonable lower-bound figure for the case fatality ratio in the general population (he was talking about the US situation) was a mere 0.05 per cent! And had they done so, the case for radical lockdowns would have disappeared, along with their opportunities for vast emergency budgets and medial glory.

So instead, it seemed that even with life and death issues of public health, the strategy for computer modellers and governments alike

would be one nicely summed up by the words of the French philosopher Jean-Jacques Rousseau centuries ago, speaking of social life in general: 'Let us begin, then, by laying facts aside, as they do not affect the question.'

All of which highlights the thinking error well known in computer programming circles, the one dubbed, GIGO – standing for 'Garbage In, Garbage Out'. However, if the problem is well known within computing, it is not well known enough outside of the field – with the result that the pronouncements of computers and computer models are accorded far more weight than anything produced by a human being – even though what the computer says is determined by the information fed into it by humans.

Tasked with solving the corona virus health crisis, the models predicted that mask mandates would slow the spread of the virus significantly, and that if you but add in other social-distancing measures and a general vaccination programme, then hey presto, you have squashed the illness flat as a pancake in just the few days it takes for the policies to take effect.

Funnily enough, though, all over the world in *real* life, the numbers of people testing positive for the virus, the numbers getting ill enough to go to doctors and hospitals, and finally the people dying if not actually from Covid, at least 'with' Covid, all went through the roof. Take these charts for Quebec, subjected to zealous pro-vaccine policies by Canadian Prime Minister Justin Trudeau. (The same Trudeau whose family, it has been alleged, has fiduciary connections to MRNA technologies through investments in Vancouver-based Acuitas Therapeutics, a company that specialises in the development of lipid nanoparticles. This is a technology that on its own is linked to heart problems[11] – something to bear in mind when allegations about footballers dropping dead from heart attacks on the pitch are being aired!)

To be fair, the charts show both masks and vaccine passports having a short 'window' of effect – keeping the daily case numbers down for a couple of months. The Quebec mask mandate came in on 18 July 2020 – northern-hemisphere summer – at a time of very low virus rates, yet by the arrival of the winter, the effect seemed to no longer hold. Vaccine passports came in on 1 September 2021, again a time in the northern hemisphere of very low circulation of respiratory 'winter' viruses, and this again seemed to work until, err… – the arrival of winter. At which point, like a gambler betting all their remaining funds, Trudeau went for a harsh lockdown. This came in on 20 December 2021, and yet cases sky-rocketed.[12] But that's dirty *real world* data – not crisp, ultra-precise computer models. And in Canada, as everywhere, policy continued to follow the computer predictions, not the evidence of real life.

You see, the great problem with computer models, the problem identified as early as 1964 by Joseph Weizenbaum, in an era when computers didn't have any of the aura that they do today, is that their pronouncements seem to have greater authority than any mere human individual can ever have.

So, when the computer models suggested, in their unchallengeable way, that so-called self-isolation plus social distancing was the universal solution to 'slowing the spread' of the virus, mere human reasoning – arguments – was, to coin a term that has lately been rather contaminated, 'trumped'. Instead, once the computer simulations appeared to identify with incredible precision how and where a virus would spread, when and where people would sicken and die, all that was left for politics was to implement the response: tackle questions like how to get the medical staff their cumbersome protective clothing; whether to support the police with soldiers in enforcing curfews and patrolling the streets; and above all, get behavioural psychologists involved in 'selling' unprecedented, legally enforced changes to the lives of the whole population.

The 'Nudge Unit' in Britain is creepy enough, but actually, much of the response was not even

designed by humans; instead, here too, it was done by computer models. As an IBM blog recently put it: 'Artificial intelligence is used to sift through the enormous amount of data necessary for contact tracing and build epidemiological models, helping contain the virus's spread.'[13]

Strange to say, though, however carefully the strategies were applied, they didn't seem to work. Indeed, the figures from many unfortunate countries were so bad that it looked, if anything, as if masks, lockdowns and even vaccination were making things worse. Cases, hospitalisations and deaths seemed to follow a standard 'bell shape' curve everywhere – low in the summer, rising through the winter and early spring.

Two years later, we actually have real data to go along with the imaginary worlds described by the modellers. So what does this data say about the efficacy of lockdowns in combatting the corona virus?

One paper in *Studies in Applied Economics* journal, by Jonas Herby, Lars Jonung and Steve H. Hanke, concluded that:[14] 'lockdowns are not an effective way of reducing mortality rates during a pandemic, at least not during the first wave of the COVID-19 pandemic.'

The study, much discussed on social media to the disgust of the mainstream media, was a meta-analysis of several studies by a Johns Hopkins University professor and others that found that lockdowns during the first Covid wave in the spring of 2020 only reduced Covid mortality by a mere 0.2 per cent in the USA and Europe. 'While this meta-analysis concludes that lockdowns have had little to no public health effects, they have imposed enormous economic and social costs where they have been adopted', the researchers wrote. 'In consequence, lockdown policies are ill-founded and should be rejected as a pandemic policy instrument.'

How could that be? The computers had spoken! Sacrilege.

The mainstream media knew what to do with such a report. CNN, MSNBC, ABC, CBS and NBC all ignored it, and concentrated instead on shaming 'red states' with minimal restrictions as 'super-spreaders'. The *New York Times,* the *Washington Post,* the *Associated Press, Reuters, USA Today, Axios* and *Politico* largely turned a blind eye to the findings or, like Bruce Lee, in *Forbes* magazine (6 February 2020), actively took up the cudgels.[15] First of all, Lee pointed out that although the study was being called a 'Johns Hopkins University' study, 'the University itself didn't write the paper, because buildings can't type on laptops without crushing them'.

We could apply that argument to every study attributed to an institution, though. Did I say that Bruce's background is… computer modelling? (And 'systems'.) Which brings me to his second objection to the study: it was by people who were not medically qualified, not even (we might suppose) 'epidemiologists'. Why would the views of economists even be sought on matters like lockdowns? For Lee, it was an absurdity. As if, he writes, 'you were to end up in the emergency room with an injury' and be told, 'Don't worry, an economist will be around shortly to re-attach your arm'. This, Lee quipped, would not be the most comforting thing to hear. He doesn't seem to realise that you might not be enormously reassured to be told that a computer expert with a mathematical model would be along in a moment, either.

No, what the mainstream media did find time for, however, was a rival study that did find lockdowns having a significant effect.

Nadya Johanna, Henrico Citrawijaya, and Grace Wangge found,[16] in contrast, that 'for lockdown, ten studies consistently showed that it successfully reduced the incidence, onward transmission, and mortality rate of COVID-19'.

How to explain the discrepancy? But first, Johanna et al. included modelling studies. Indeed, out of a total of 14 studies, ten were not real-life observations, but computer models. An additional cherry was that, 'based on one modelling' study, mass testing reduced the total infected people compared to no mass testing.

These authors made no grand claims to respected institutions; rather, they were modest figures hailing from modest Indonesian institutes. Johanna, 'a community organiser, and a public speaker, focusing on health promotion and prevention', claims in addition to have treated some Covid patients as a general practitioner. Clearly, this is a case where the research was the news, not the researchers. Research which simply amplified the results of computer modellers. In programming, this is what's know as a feedback loop. And such things are very dangerous.

But guess which study was adopted by governments as representing 'the science'? If members of the public think that 'real life' is more important than computer models, actually, the experts think the gold standard is what happens in the model. And so it is that for the likes of Aleksandra Mojsilović, an IBM Fellow and head of Artificial Intelligence Foundations at IBM Research, who runs a centre ominously called 'Science for Social Good', the models are correct, but real life is messy. Here, the computer modellers are modelling why the models don't work. And the built-in assumption in these models is that the problem is with real life.

In an interview for the IBM blog,[17] she says researchers are already beginning to think about how to process the multitude of decisions made to determine what worked and what didn't, adding:

> All the events that we've collected so far in terms of the kinds of decisions that people made during the epidemic – government decisions, school closures, decisions to wear a mask or not – machine learning and AI are giving us toolkits to essentially analyse all of this data.

'The Centre for Social Good' is full of mathematicians and computer programmers and, of course, no social scientists, anthropologists or philosophers. George Orwell could explain that. One curious result is that when they looked at the qualities that made 'applicants' successful in what they call 'social enterprises', they did so by correlating the use of certain key words in applications with later promotions. They found that 'talking about "impact" and "team work" makes applicants more likely to qualify through the evaluation phases', and that applicants who identified 'climate change as a problem area', and green technology and energy as part of the solution, were exactly the right people to hire. This is all so silly it would be laughable but, believe me, it isn't. The people with the computers have the money and the influence today.

Actually, as far as I know, the IBM labs are not particularly implicated in lockdown policies, although they certainly are involved with things like the modelling of new vaccines. (IBM's baffling new technology of query-based molecular optimisation (QMO), is an AI framework that can help improve discovery workflows and accelerate the delivery of new molecules and materials.)

Perhaps the vaccines are the real standard-bearers of computing in this virus saga. Because, for sure, this is a computer-designed vaccine. Dr Jason Crain, a researcher at IBM and visiting professor at the University of Oxford, also spoke of there being 'eye-watering amounts of computational power being directed to the coronavirus challenge' in the blog (for IBM) mentioned above.

'There are about 8 billion people on Earth at the moment and if you could imagine each of them doing a million calculations per person in the world every second, that would be something like eight petaflops of computing power', says Professor Crain. 'The COVID-19 High

Performance Computing Consortium (HPC), supercomputers at the world's largest tech firms and advanced laboratories, including NASA, Los Alamos and Berkeley, have made available something like 50 times that.'

So how do they use all this computing power? Well, corona viruses use 'spikes' to penetrate and then infect cells and so, under traditional techniques, researchers would try to grow the virus in a laboratory, and then introduce new compounds to observe which ones bind to the spike to render it ineffective.

What a useless method! No, nowadays, computer simulation has replaced that process, and generates millions, or even billions, of pieces of data in the blink of an eye. In fact, they generate so much data that standard research computers cannot easily cope with it.

Harnessing this fabulous global array of computer power, however, researchers 'were able to simulate 8000 compounds in a matter of days, identifying 77 small molecules' that the models showed had 'the potential to inhibit COVID-19'.

Did the models look at all the possible other 'side effects' of these 'compounds'? Of course not. You'd need a computer model for the whole complexity of the human body; that's not been developed yet. Indeed (hint to computer developers), there isn't one standard program running in humans – our immune systems are not only fabulously complex, well beyond the computer models, but to a significant extent unique to each of us.

However, what I really want to emphasise in this article is how incredibly grand, and even intimidating, today's computer technology is. It's easy to think that only computer experts, maybe even only other computers, can possibly challenge it. However, to think this would be wrong. Just like the theory of relativity, the principles underlying computer models, and the assumptions built into them, both can and *need to be* challenged.

Today, in 2022, the corona virus provides a case study of how dodgy assumptions and duff statistics can be laundered through grand computer programs, creating in the process a Ten Trillion Dollar 'pandemic'[18] with enormous social costs. Maybe common sense says that this is a pandemic that perhaps never was. But increasingly we don't refer to common sense for our view of the world; we refer to computer simulations.

Notes and References

1 https://tinyurl.com/27mjntnm. See also Peter Aitken, 'Imperial College model Britain used to justify lockdown a "buggy mess", "totally unreliable", experts claim', Fox News, 16 May 2020; available at https://tinyurl.com/2hkvcajj (accessed 20 May 2022).
2 Martin Cohen, *Paradigm Shift: How Expert Opinions Keep Changing on Life, the Universe, and Everything*, Imprint Academic, Exeter. 2015.
3 See James Sturcke, 'Burning issue', the *Guardian*, 22 August 2005; available at https://tinyurl.com/bdhuf8et (accessed 20 May 2022).
4 Maggie Koerth, Laura Bronner & Jasmine Mithani, 'Why it's so freaking hard to make a good COVID-19 model', 31 March 2020; available at https://tinyurl.com/2jd77tnw (accessed 20 May 2022).
5 Kevin Dayaratna and Norbert Michel describe how and why here: 'Failures of an influential COVID-19 model used to justify lockdowns', The Heritage Foundation, 18 May 2020; available at https://tinyurl.com/2574cfbk (accessed 20 May 2022).
6 Link: https://tinyurl.com/3a5wd5mu – not available at 20 May 2022; see, instead, https://tinyurl.com/j89vu2em (accessed 20 May 2022).
7 See Berkeley Lovelace Jr and Noah Higgins-Dunn, 'WHO says coronavirus death rate is 3.4% globally, higher than previously thought', CNBC, 3 March 2020; available at https://tinyurl.com/2p8eu78m (accessed 20 May 2022).
8 See John P.A. Ioannidis, 'A fiasco in the making? As the coronavirus pandemic takes hold, we are making decisions without reliable data', *STAT* magazine, 17 March 2020; available at

https://tinyurl.com/yc3ska43 (accessed 20 May 2022).

9 See Sabine L. van Elsland, 'COVID-19 deaths: infection fatality ratio is about 1% says new report', Imperial College London News, 29 October 2020; available at https://tinyurl.com/5ak6n8f7 (accessed 20 May 2022).

10 See The Great Barrington Declaration, 4 October 2020; available at https://gbdeclaration.org (accessed 20 May 2022).

11 See Nanotechnologies, Level 2, available at https://tinyurl.com/mr38f95x (accessed 20 May 2022).

12 See Daily new cases per million – Quebec; available at https://tinyurl.com/55mhusjf (accessed 20 May 2022). Original source, Johns Hopkins University CSSE Github.

13 See Anon, 'Supercomputers: Supercomputing a COVID-19 vaccine', IMB Think Broadcast, 23 July 2020; available at https://tinyurl.com/59y4wwvt (accessed 20 May 2022).

14 See Jonas Herby, Lars Jonung & Steve H. Hanke, 'A literature review and meta-analysis of the effects of lockdowns on COVID-19 mortality', *Studies in Applied Economics*, 200, Jan. 2022; available at https://tinyurl.com/3345kcjf (accessed 20 May 2022).

15 See Bruce Y. Lee, 'Did so-called "Johns Hopkins Study" really show lockdowns were ineffective against Covid-19?', Forbes, 6 February 2022; available at https://tinyurl.com/4rn4ckcf (accessed 20 May 2022).

16 See Nadya Johanna, Henrico Citrawijaya & Grace Wangge, 'Mass screening vs lockdown vs combination of both to control COVID-19: a systematic review', *Journal of Public Health Research*, 9 (4), 14 October 2020, pp. 523–31; available at https://tinyurl.com/4ymmj9d7 (accessed 20 May 2022).

17 See note 13.

18 See Frank Holmes, 'A $10 trillion response to the global pandemic', Forbes, 23 March 2020; available at https://tinyurl.com/r66jxfbe (accessed 20 May 2022)

About the contributor

Dr Martin Cohen is an author specialising in popular books in philosophy and social science, who likes making connections between different areas and ideas. His latest book *Rethinking Thinking: Problem Solving from Sun Tzu to Google*, was published by Imprint in April. As well as his previously featured book *Paradigm Shift*, his other books include *101 Philosophy Problems*, *Critical Thinking Skills for Dummies*, *I Think Therefore I Eat* and even a book on nuclear economics called *The Doomsday Machine*. *I Think Therefore I Eat* combines philosophy and food, offering surprising insights into why everything we eat makes us fat, and seems to have more to do with laboratories than farms!

SOME HUMANISTIC WISDOM

"If you only have a hammer, you tend to see every problem as a nail."

Abraham Maslow (1908–1970)